arXiv:1403.1056v1 [cs.CV] 5 Mar 2014

K-TANGENT SPACES ON RIEMANNIAN MANIFOLDS FOR IMPROVED PEDESTRIAN DETECTION

Andres Sanin, Conrad Sanderson, Mehrtash T. Harandi, Brian C. Lovell

NICTA, PO Box 6020, St Lucia, QLD 4067, Australia
University of Queensland, School of ITEE, QLD 4072, Australia

ABSTRACT

For covariance-based image descriptors, taking into account the curvature of the corresponding feature space has been shown to improve discrimination performance. This is often done through representing the descriptors as points on Riemannian manifolds, with the discrimination accomplished on a tangent space. However, such treatment is restrictive as distances between arbitrary points on the tangent space do not represent true geodesic distances, and hence do not represent the manifold structure accurately. In this paper we propose a general discriminative model based on the combination of several tangent spaces, in order to preserve more details of the structure. The model can be used as a weak learner in a boosting-based pedestrian detection framework. Experiments on the challenging INRIA and DaimlerChrysler datasets show that the proposed model leads to considerably higher performance than methods based on histograms of oriented gradients as well as previous Riemannian-based techniques.

Index Terms— pedestrian detection, covariance descriptors, Riemannian manifolds, tangent spaces, boosting.

1. INTRODUCTION

Covariance descriptors have recently become popular features for pedestrian detection, due to their relative low sensitivity to translations and scale variations, as well as being quick to compute through integral image representations [5, 11, 14]. Covariance descriptors belong to the group of positive definite symmetric matrices (Sym_d^+), which can be interpreted as points on Riemannian manifolds [7]. Intuitively, we can think of a manifold as a continuous surface lying in a higher dimensional Euclidean space. By exploiting the curvature of the feature space formed by such matrices, higher performance can be achieved than vector space approaches such as histograms of oriented gradients (HOGs), previously considered as the state of the art features for pedestrian detection [2].

While the geometry of Riemannian manifolds has been extensively studied, with well defined metrics and operations for Sym_d^+, classification directly on manifolds (eg. dividing a manifold into two distinct parts) is still a challenging problem [8, 13]. Current approaches typically first map the points on the manifold to a tangent space, where traditional machine learning techniques can be used for classification [6, 14]. A tangent space is an Euclidean space relative to a point in the manifold (the tangent pole). Processing a manifold through a single tangent space is necessarily restrictive, as only distances to the tangent pole are true geodesic distances. Distances between arbitrary points on the tangent space do not represent true geodesic distances, and hence do not represent the manifold structure accurately.

As a partial workaround to the above limitation, Tuzel et al. [14] recently proposed a pedestrian detection method based on a LogitBoost cascade [4] over Riemannian manifolds. Linear regression was used as a weak classifier within the cascade, after mapping the selected training points to a tangent space that occurs at the mean of the points. In other words, for each weak classifier a new tangent pole is obtained closer to the most critical samples. As a consequence of boosting, the final inference on the manifold was made through multiple tangent spaces. However, there is an inherent limitation in that each boosting iteration has to deal with two problems at the same time: finding the best tangent representation and finding the best discriminative model.

In this paper, we propose a general discriminative model based on the combination of several tangent spaces, in order to preserve more details of the manifold structure. We shall refer to this model as *k*-tangent spaces. The model can be used as a weak classifier in a boosting-based framework, allowing the learning algorithm to focus on better classification rather than finding the best representation of the manifold. We demonstrate empirically that the proposed model leads to considerably improved pedestrian detection performance.

We continue the paper as follows. In Section 2 we overview relevant manifold operations for the space of symmetric positive definite matrices. The proposed model is described in Section 3. Results from pedestrian detection experiments on two challenging datasets are given in Section 4. The main findings and possible future directions are summarised in Section 5.

2. RIEMANNIAN GEOMETRY

In this section we briefly review Riemannian geometry, with a focus on the space of symmetric positive definite matrices. Formally, a manifold is a topological space which is locally similar to an Euclidean space [7, 14]. The minimum length curve connecting two points on the manifold is called the geodesic, and the distance between two points $\boldsymbol{X}$ and $\boldsymbol{Y}$ is given by the length of this curve.

The tangent space, $T_{\boldsymbol{X}}$ at $\boldsymbol{X}$, is the plane tangent to the surface of the manifold at that point. Geodesics (on the man-

Acknowledgements: NICTA is funded by the Australian Government as represented by the *Department of Broadband, Communications and the Digital Economy*, as well as the Australian Research Council through the *ICT Centre of Excellence* program.

ifold) are related to the vectors in the tangent space. Two operators, namely the exponential map $\exp_{\boldsymbol{X}}$ and the logarithm map $\log_{\boldsymbol{X}} = \exp_{\boldsymbol{X}}^{-1}$, switch between manifold and tangent space at $\boldsymbol{X}$. For Sym_d^+ the exponential and logarithm maps are defined as:

$$\exp_{\boldsymbol{X}}(\boldsymbol{y}) = \boldsymbol{X}^{\frac{1}{2}} \exp\left(\boldsymbol{X}^{-\frac{1}{2}} \boldsymbol{y} \boldsymbol{X}^{-\frac{1}{2}}\right) \boldsymbol{X}^{\frac{1}{2}} \quad (1)$$

$$\log_{\boldsymbol{X}}(\boldsymbol{Y}) = \boldsymbol{X}^{\frac{1}{2}} \log\left(\boldsymbol{X}^{-\frac{1}{2}} \boldsymbol{Y} \boldsymbol{X}^{-\frac{1}{2}}\right) \boldsymbol{X}^{\frac{1}{2}} \quad (2)$$

where $\boldsymbol{y}$ (on the tangent space) is a representative of $\boldsymbol{Y}$ (on the manifold). Furthermore, $\exp(\cdot)$ and $\log(\cdot)$ are matrix exponential and logarithm operators, respectively. For symmetric positive definite matrices they can be computed through Singular Value Decomposition (SVD). More specifically, let $\boldsymbol{X} = \boldsymbol{U}\Sigma\boldsymbol{U}^T$ be the SVD of the symmetric matrix $\boldsymbol{X}$, then

$$\exp(\boldsymbol{X}) = \boldsymbol{U} \exp(\Sigma) \boldsymbol{U}^T \quad (3)$$

$$\log(\boldsymbol{X}) = \boldsymbol{U} \log(\Sigma) \boldsymbol{U}^T \quad (4)$$

where $\exp(\Sigma)$ and $\log(\Sigma)$ are diagonal matrices with diagonal elements that are respectively equivalent to the exponential or logarithms of the diagonal elements of matrix Σ.

Using the tangent space defined at $\boldsymbol{X}$, the distance between $\boldsymbol{X}$ and $\boldsymbol{Y}$ can be calculated via:

$$d^2(\boldsymbol{X}, \boldsymbol{Y}) = \text{trace}\left\{\log^2\left(\boldsymbol{X}^{-\frac{1}{2}} \boldsymbol{Y} \boldsymbol{X}^{-\frac{1}{2}}\right)\right\} \quad (5)$$

Given a set $\{\boldsymbol{X}_i\}_{i=1}^N$ of points on a manifold $\mathcal{M}$, the Karcher mean [9] is the point on $\mathcal{M}$ that minimises the sum of distances:

$$\boldsymbol{\mu} = \arg\min_{X \in \mathcal{M}} \sum\nolimits_{i=1}^{N} d^2(\boldsymbol{X}_i, \boldsymbol{X}) \quad (6)$$

For Sym_d^+, the Karcher mean can be iteratively found using:

$$\boldsymbol{\mu}^{[t+1]} = \exp_{\boldsymbol{\mu}^{[t]}}\left\{N^{-1} \sum\nolimits_{i=1}^{N} \log_{\boldsymbol{\mu}^{[t]}}(\boldsymbol{X}_i)\right\} \quad (7)$$

3. K-TANGENT SPACES FOR PEDESTRIAN DETECTION

Traditional tangent-based discrimination approaches generally have three steps: **(i)** select the tangent pole $\boldsymbol{X} \in \mathcal{M}$; **(ii)** project the training samples onto the tangent space $T_{\boldsymbol{X}}$; and **(iii)** train a discriminative model g (eg. the regression function in [14]) using the projected samples. The discrimination performance hence depends on two characteristics: **(a)** how well model g is able to fit the data, and **(b)** how well the tangent representation preserves the distances between the points on the manifold. Intuitively speaking, the larger the area covered by the tangent space, the worse the representation.

To address the above issue, we propose the combination of multiple tangent spaces into a single discriminative model:

$$\lambda = \{(\boldsymbol{\mu}_k, n_k, g_k)\}_{k=1}^K \quad (8)$$

where, for the k-th entry, $\boldsymbol{\mu}_k$ is the tangent pole (center), n_k is the number of samples associated with the center, and g_k is a discriminative model associated with the center. By having several spaces, the likelihood of each sample point having at least one appropriate vector space representation is increased.

Fig. 1: Examples of pedestrian images in **(a)** INRIA and **(b)** DaimlerChrysler datasets.

For a given test sample $\boldsymbol{Z} \in \mathcal{M}$, the combined output of the K discriminative models can be obtained using:

$$G(\boldsymbol{Z}|\lambda) = \sum_{k=1}^{K} \frac{n_k}{\sum_{l=1}^{K} n_l} g_k\left(\log_{\boldsymbol{\mu}_k}(\boldsymbol{Z})\right) \quad (9)$$

where each n_k is used as a mixing coefficient, emphasising more dense clusters that are less likely to be affected by outliers.

Within the context of pedestrian detection, the proposed model is trained by first grouping positive samples into K clusters. Negative samples are not taken into account at this stage as they are not well characterised (ie. they do not represent a coherent class). To obtain the cluster centers, a form of the k-means algorithm [3] explicitly adapted to manifolds can be used. Specifically, the geodesic distance in Eqn. (5) is used as the distance measure, with each cluster center taken to be a Karcher mean, as in Eqn. (6). Once the cluster centers are obtained, all samples are projected onto each tangent space k and a discriminative model g_k is trained. The process is summarised in Algorithm 1.

The *k*-tangent spaces model can be used to replace the weak classifier (which uses a single tangent space) within the boosting-based algorithm proposed by Tuzel et al. [14]. The discriminative models within each *k*-tangent spaces model are the same as the original weak learners (regression functions) used in [14]. The resulting strong classifier produced by boosting is hence comprised of a cascade of *k*-tangent spaces.

Algorithm 1 Training the *k*-tangent spaces model for pedestrian detection

Input:

- Training samples and their labels: $\{(\boldsymbol{X}_i, y_i)\}_{i=1}^N$, $\boldsymbol{X}_i \in \mathcal{M}, y_i \in \{-1, +1\}$

Processing:

- Cluster positive samples to obtain K cluster centers and the number of samples associated with each center: $\{(\boldsymbol{\mu}_k, n_k)\}_{k=1}^K$
- For $k = 1, \ldots, K$:
  - Map all samples to the tangent space at $\boldsymbol{\mu}_k$ via $\log_{\boldsymbol{\mu}_k}(\cdot)$
  - Train discriminative model g_k using the mapped samples and their labels

Output:

- Model $\lambda = \{(\boldsymbol{\mu}_k, n_k, g_k)\}_{k=1}^K$

4. EXPERIMENTS

We used two datasets of pedestrian images for the experiments: INRIA [1] and DaimlerChrysler [10]. Examples are shown in Fig. 1.

The INRIA dataset contains 2416 pedestrian samples and 1218 person-free images for training, as well as 1126 pedestrian samples and 453 person-free images for testing. The size of the samples is 64×128 with a margin of 16 pixels around the pedestrians. The main challenges of the INRIA dataset are variations in pose, background and partial occlusions.

The DaimlerChrysler dataset contains three training sets and two tests sets. Each training and test set contains 4800 pedestrian samples and 5000 negative samples. Additionally, each training set contains 1200 person-free images where further negative samples can be obtained. The size of the positive samples is 18×36 with a margin of 2 pixels around the pedestrians. Detection on this dataset is more difficult due to the small size of the samples and the challenging negative set.

We performed three experiments: **(1)** exploration of the effect of the number of tangent spaces on detection performance; **(2)** comparison of the detection performance against other tangent mapping approaches; **(3)** comparison of the performance against several notable pedestrian detection methods in the literature. In all cases the proposed *k*-tangent spaces model was used as a weak learner within the boosting-based pedestrian detection framework proposed by Tuzel et al. [14].

As per [14], each pixel $I_{(x,y)}$ located at (x, y) was represented by an 8-dimensional feature vector:

$$\boldsymbol{f}_{(x,y)} = \left[x \;\; y \;\; |d_x| \;\; |d_y| \;\; |d_{xx}| \;\; |d_{yy}| \;\; (d_x^2 + d_y^2)^{\frac{1}{2}} \;\; \arctan(|d_x|/|d_y|) \right]$$

where $d_x = \partial I_{(x,y)}/\partial x$, $d_{xx} = \partial^2 I_{(x,y)}/\partial x^2$ with the last two terms representing edge magnitude and orientation. The covariance descriptor of a given region of pixels (selected by the boosting framework) is hence an 8×8 matrix.

In the first experiment we varied the number of tangent spaces from 1 to 5. To avoid training on the background in the INRIA dataset, the margin around the pedestrians in the positive samples was reduced from 16 to 2 pixels. The results are presented as detection error trade-off curves in Fig. 2, where the best performance point is at the bottom-left corner the graph (ie. minimal error rates). The results show that the detection performance generally increases as the number of spaces increases to 4. Using 5 tangent spaces degrades performance. We conjecture that the number of training samples affects the optimal number of spaces. The more spaces are used, the lower the number of samples is available for each cluster, and hence it is more likely that the clusters are affected by outliers.

In the second experiment, we compared the performance of four tangent mapping approaches: **(i)** directly converting the covariance descriptors into vectors (ie. without projection to a tangent space); **(ii)** projecting the points to the tangent space with the identity matrix as the pole; **(iii)** projecting the points to the tangent space with the Karcher mean of positive samples as the pole; **(iv)** using the proposed *k*-tangent spaces model with 3 spaces (time limitations prevented us from using 4 spaces).

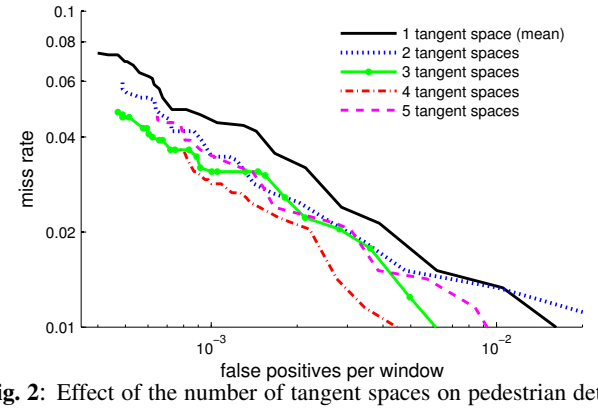

Fig. 2: Effect of the number of tangent spaces on pedestrian detection performance on the INRIA dataset.

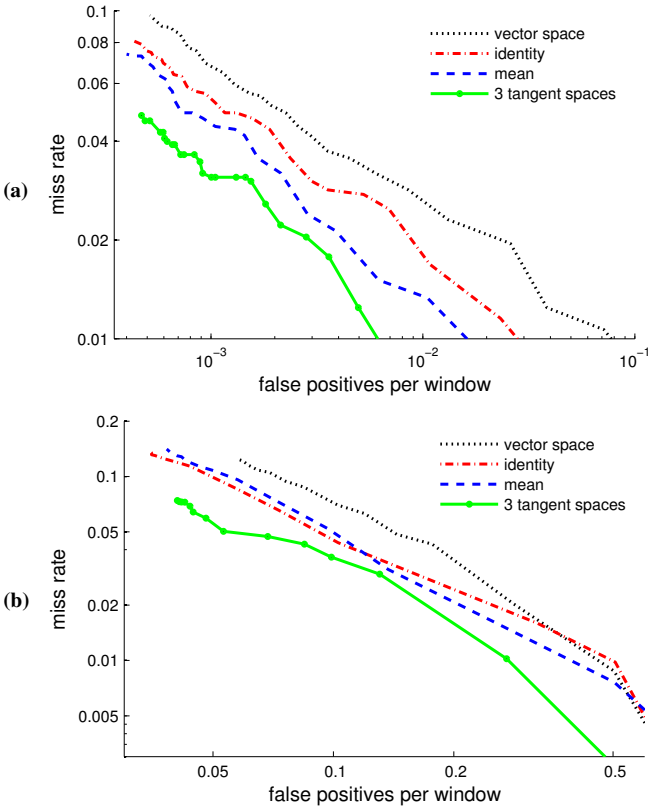

Fig. 3: Pedestrian detection using various tangent space representations on **(a)** INRIA and **(b)** DaimlerChrysler.

The results, presented in Fig. 3, indicate that any tangent space representation leads to better performance compared to the direct vector space representation. On the INRIA dataset, the Karcher mean mapping outperforms the identity matrix mapping, in agreement with the results presented in [14]. However, on the more challenging DaimlerChrysler dataset, there is little difference between the performance of the Karcher mean and identity matrix mapping. The proposed *k*-tangent spaces model leads to the best performance on both datasets.

In the third experiment we compared the proposed approach against several notable pedestrian detection methods in the literature. For the INRIA dataset, we selected

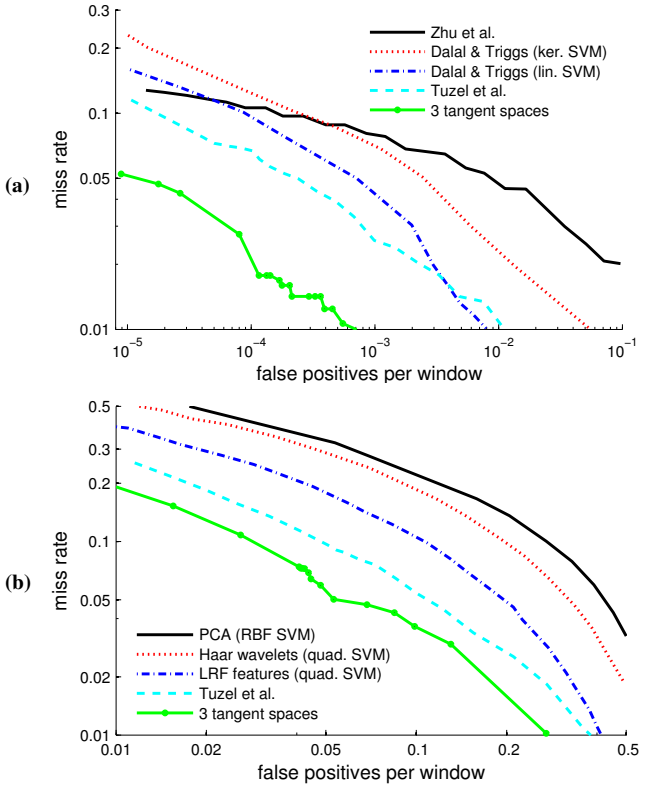

Fig. 4: Comparison with other pedestrian detection methods on **(a)** INRIA and **(b)** DaimlerChrysler. The configuration of pedestrian images was different than used in Fig. 3, to allow direct comparison with previously published results [1, 10, 14, 15].

three techniques based on HOG features: two SVM based methods [1], and an AdaBoost based method [15]. For the DaimlerChrysler dataset, we selected the best results for the three descriptors proposed in [10]: PCA coefficients, Haar wavelets, and local receptive fields. For both INRIA and DaimlerChrysler, we have also compared against the method of Tuzel et al. [14], where mapping based on the Karcher mean was used (ie. single tangent space).

To allow a direct comparison with previously published results, the original images in the INRIA dataset were used (ie. with a margin of 16 pixels around the pedestrians), rather the margin reduced images as used in the first and second experiments. The results, presented in Fig. 4, indicate that on both datasets considerably better detection performance is obtained by the proposed k-tangent spaces approach.

5. MAIN FINDINGS AND FUTURE DIRECTIONS

Covariance descriptors have been recently demonstrated to be useful features for pedestrian detection, due to their relative low sensitivity to translations and scale variations. By interpreting the descriptors as points on Riemannian manifolds, and by exploiting the curvature of such manifolds, higher performance can be achieved than vector space approaches such as histograms of oriented gradients.

Current manifold-based classification approaches typically map the points to a tangent space, where traditional machine learning techniques are used. However, such treatment is restrictive as distances between arbitrary points on the tangent space do not represent true geodesic distances, and hence do not represent the manifold structure accurately.

In this paper we proposed a general discriminative model based on the combination of several tangent spaces, in order to preserve more details of the structure. By having several spaces, the likelihood of a sample point having at least one appropriate vector space representation is increased. The model was used as a weak learner within the boosting-based pedestrian detection framework proposed by Tuzel et al. [14]. Experiments on the challenging INRIA and DaimlerChrysler datasets indicate that the proposed model (with just a few tangent spaces) leads to considerably improved performance.

Future work includes adapting the model to more general classification problems (ie. not tied to covariance matrices or pedestrian detection [12]), as long as the samples can be represented as points on Riemannian manifolds.

6. REFERENCES

[1] N. Dalal and B. Triggs. Histograms of oriented gradients for human detection. In *IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 886–893, 2005.

[2] P. Dollar, C. Wojek, B. Schiele, and P. Perona. Pedestrian detection: an evaluation of the state of the art. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 34(4):743–761, 2012.

[3] C. Elkan. Using the triangle inequality to accelerate k-means. In *International Conference on Machine Learning*, pages 147–153, 2003.

[4] J. Friedman et al. Additive logistic regression: a statistical view of boosting. *Annals of Statistics*, 28(2):337–407, 2000.

[5] G. Gualdi et al. Contextual information and covariance descriptors for people surveillance: an application for safety of construction workers. *EURASIP Journal on Image and Video Processing*, 2011.

[6] K. Guo et al. Action recognition using sparse representation on covariance manifolds of optical flow. In *IEEE Int. Conf. Advanced Video and Signal Based Surveillance (AVSS)*, pages 188–195, 2010.

[7] M. T. Harandi, C. Sanderson, R. Hartley, and B. C. Lovell. Sparse coding and dictionary learning for symmetric positive definite matrices: a kernel approach. In *European Conference on Computer Vision (ECCV), Lecture Notes in Computer Science (LNCS)*, volume 7573, pages 216–229, 2012.

[8] M. T. Harandi, C. Sanderson, A. Wiliem, and B. C. Lovell. Kernel analysis over Riemannian manifolds for visual recognition of actions, pedestrians and textures. In *IEEE Workshop on the Applications of Computer Vision*, pages 433–439, 2012.

[9] H. Karcher. Riemannian center of mass and mollifier smoothing. *Comm. Pure and Applied Math.*, 30:509–541, 1977.

[10] S. Munder and D. Gavrila. An experimental study on pedestrian classification. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 28(11):1863–1868, 2006.

[11] S. Paisitkriangkrai, C. Shen, and J. Zhang. Fast pedestrian detection using a cascade of boosted covariance features. *IEEE Trans. Circuits and Systems for Video Technology*, 18(8):1140–1151, 2008.

[12] C. Sanderson, M. T. Harandi, Y. Wong, and B. C. Lovell. Combined learning of salient local descriptors and distance metrics for image set face verification. In *IEEE International Conference on Advanced Video and Signal-Based Surveillance (AVSS)*, pages 294–299, 2012.

[13] S. Shirazi, M. T. Harandi, C. Sanderson, and B. C. Lovell. Clustering on Grassmann manifolds via kernel embedding with application to action analysis. In *IEEE International Conference on Image Processing*, pages 781–784, 2012.

[14] O. Tuzel, F. Porikli, and P. Meer. Pedestrian detection via classification on Riemannian manifolds. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 30(10):1713–1727, 2008.

[15] Q. Zhu et al. Fast human detection using a cascade of histograms of oriented gradients. In *IEEE Conf. Computer Vision and Pattern Recognition*, pages 1491–1498, 2006.